

Culture, embodiment and genes: unravelling the triple helix

Michael Wheeler^{1,*} and Andy Clark²

¹*Department of Philosophy, University of Stirling, Stirling FK9 4LA, UK*

²*School of Philosophy, Psychology, and Language Sciences, David Hume Tower, George Square, Edinburgh EH8 9J7, UK*

Much recent work stresses the role of embodiment and action in thought and reason, and celebrates the power of transmitted cultural and environmental structures to transform the problem-solving activity required of individual brains. By apparent contrast, much work in evolutionary psychology has stressed the selective fit of the biological brain to an ancestral environment of evolutionary adaptedness, with an attendant stress upon the limitations and cognitive biases that result. On the face of it, this suggests either a tension or, at least, a mismatch, with the symbiotic dyad of cultural evolution and embodied cognition. In what follows, we explore this mismatch by focusing on three key ideas: cognitive niche construction; cognitive modularity; and the existence (or otherwise) of an evolved universal human nature. An appreciation of the power and scope of the first, combined with consequently more nuanced visions of the latter two, allow us to begin to glimpse a much richer vision of the combined interactive potency of biological and cultural evolution for active, embodied agents.

Keywords: cultural transmission; embodied cognition; niche construction; evolutionary psychology; modularity; neuroconstructivism

1. INTRODUCTION: A TENSION REVEALED

There is a natural affinity between work that stresses the role of embodiment and action in thought and reason (examples include Varela *et al.* 1991; Clark 1997; Noë 2004; Wheeler 2005) and work that explores the cognitive role of cultural evolution (Tomasello 1999; Kirby 2002; Sterelny 2003). Both approaches share an emphasis on the power of non-neural structures to transform the shape of the problem-solving activity required of individual brains. Such potent non-neural structures take a wide variety of forms, from the biomechanics of the gross physical body (Collins *et al.* 2005), to the structural features of a linguistic code (Kirby 2002), and on to aspects of the local, physical and social environment (for some reviews, see Clark 1997; Wilson & Clark *in press*). Many of these enabling non-neural structures are self- or species created, and are thus both products and determinants of human thought and activity. Such products and determinants are also subject to cycles of transmission, alteration and inheritance, in at least a rough analogy with genetic inheritance systems (e.g. Jablonka & Lamb 2005). The result (as we shall see) is a vision of the evolution, the development and the real-time unfolding of human cognition, in which a kaleidoscope of complex ratchet effects fuel the flexible and, to a significant degree, open-ended character of thought and action.

By apparent contrast, much work in evolutionary psychology¹ has stressed the selective fit of the biological brain to some ancestral environment of evolutionary adaptedness, with an attendant focus

upon the limitations and cognitive biases that result (see, canonically, Barkow *et al.* 1992. For more recent coverage, see Buss 2005). On the face of it, this suggests either a tension or, at least, a mismatch with our symbiotic dyad of cultural evolution and embodied cognition. In place of a dynamic and transformative interplay of neural, bodily and (sometimes self-created) environmental resources over different time scales, we confront a restricted set of pre-specified adapted functions, performed in the triggering context of variable non-neural structures and cultural forces, by relatively static, genetically based forms of neural encoding and processing.

In what follows, we explore this mismatch by focusing on three key ideas: cognitive niche construction; cognitive modularity; and the existence of an evolved human nature. An appreciation of the power and scope of the first, combined with consequently more nuanced visions of the latter two, allow us (we shall argue) to begin to glimpse a much richer vision of the combined potency of biological and cultural evolution for active, embodied agents. In §2, we explain the basic idea of cognitive niche construction. In §§3–5, we explore that idea in a variety of settings. The outcome is a clearer understanding of how cultural transmission and embodied cognition generate the first image of human cognitive systems identified above. That done, §§6 and 7 unpack the alternative (evolutionary-psychological) picture by focusing on the interlocking notions of cognitive modularity and an evolved human nature. In §§8–11, we endeavour to resolve some of the tension between our two visions, by examining how, and to what extent, the notions of cognitive modularity and an evolved human nature may be reconstructed within a cognitive niche-construction framework. This brings

* Author for correspondence (m.w.wheeler@stir.ac.uk).

One contribution of 11 to a Theme Issue 'Cultural transmission and the evolution of human behaviour'.

into focus what we, adapting the original usage by Lewontin (2000), are dubbing *triple helix* models of mind and cognition. These are models in which the goal is to take seriously, and ultimately to understand, the multiple ways in which three tangled sets of factors—*culture, embodiment and genes*—combine to make us the beings that we are.²

2. COGNITIVE NICHE CONSTRUCTION

Niche construction, as defined by Laland *et al.* (2000, p. 131), refers to:

the activities, choices and metabolic processes of organisms, through which they define, choose, modify and partly create their own niches. For instance, to varying degrees, organisms choose their own habitats, mates, and resources and construct important components of their local environments such as nests, holes, burrows, paths, webs, dams, and chemical environments.

Niche construction is a pervasive, though still widely underestimated, force in nature. All animals act on their environments and, in so doing, alter those environments in ways that may sometimes change the fitness landscape of the animal itself. A classic example³ is the spider's web. The existence of the web modifies the sources of natural selection within the spider's selective niche, allowing (for example) subsequent selection for web-based forms of camouflage and communication.

Still further complexity is introduced when organisms collectively build structures that persist beyond their own lifetime. A familiar example is the communally constructed beaver's dam, whose physical presence subsequently alters selection pressures on both the beaver and its progeny, who inherit the dam and the altered river flows it has produced. Similar effects can be seen in the nest-building activities of many wasps and termites, where the presence of the nest introduces selection pressures for behaviours that regulate nest temperature by (for example) sealing entrances at night (von Frisch 1975).

The cultural transmission of knowledge and practices resulting from individual lifetime learning, when combined with the physical persistence of artefacts, yields yet another source of potentially selection-impacting feedback. The classic example here (from Feldman & Cavalli Sforza 1989) is the practice of domesticating cattle and dairying, which paved the way for selection for adult lactose tolerance in (and only in) those human populations engaging in such activities.

In all these cases, what ultimately matters, as Laland *et al.* (2000) stress, is the way niche-construction activity leads to new feedback cycles. In the standard cases, these feedback cycles run across evolutionary time. Animals change the world in ways that change the selective landscapes for biological evolution. But it is worth pointing out that this whole process has a direct analogue within lifetime learning. Here, the feedback cycles alter and transform processes of individual and cultural reasoning and learning. For example, both educational practices and human-built structures (artefacts) are passed on from generation to generation in ways that dramatically alter the fitness landscape for individual lifetime learning. To adapt an example one of us has

used elsewhere (Clark 2001), the novice bartender inherits an array of differently shaped glassware and cocktail furniture, and a culturally transmitted practice of serving different drinks in different kinds of glass. As a result, expert bartenders learn to line up differently shaped glasses in spatial sequence corresponding to the temporal sequence of drinks orders (Beach 1988). The problem of remembering what drink to prepare next is thus transformed, as a result of learning within this pre-structured niche, into the problem of perceiving the different shapes and associating each shape with a kind of drink. The bartender, by creating persisting spatially arrayed stand-ins for the drinks orders, actively structures the local environment so as to press more usefulness from the basic modes of visually cued action and recall. In this way, the exploitation of the physical situation allows relatively lightweight cognitive strategies to reap large rewards.

This is a simple illustration of the power of cognitive niche construction, defined as the process by which animals build physical structures that transform problem spaces in ways that aid (or sometimes impede) thinking and reasoning about some target domain or domains.⁴ These physical structures combine with appropriate culturally transmitted practices to transform problem solving, and (in the most dramatic cases) to make possible whole new forms of thought and reason.⁵ Sections 3–5 of this paper explore the idea of cognitive niche construction in a variety of settings.

3. THINKING SPACE

A vast amount of contemporary human cognitive niche construction involves the active exploitation of space, often by way of culturally inherited artefacts and culturally transmitted strategies. Kirsh (1995) in his classic treatment 'The Intelligent Use of Space' divides these uses into three broad (and overlapping) categories. The first is 'spatial arrangements that simplify choice', such as laying out cooking ingredients in the order you will need them, or putting your shopping in one bag and mine in another. The second is 'spatial arrangements that simplify perception', such as putting the washed mushrooms on the right of the chopping board and the unwashed ones on the left, or the colour green dominated jigsaw puzzle pieces in one pile and the red dominated ones in another. The third is 'spatial dynamics that simplify internal computation', such as repeatedly reordering the scrabble pieces so as to prompt better recall of candidate words, or the use of instruments such as slide rules, which transform arithmetical operations into perceptual alignment activities.

It is noteworthy that the majority of these spatial arrangement ploys work, as Kirsh himself notes at the end of his treatment, by *reducing the descriptive complexity of the environment*. Space is often used as a resource for grouping items into equivalence classes for some purpose (e.g. *washed* mushrooms, *red* jigsaw pieces, *my* shopping and so on). Human language, perhaps the ultimate cognitive tool (Clark 1997), is itself notable for both its open-ended expressive power and its ability to reduce the descriptive complexity of the environment. Reduction of descriptive complexity, however achieved,

makes new groupings available for thought and action. In this way, the intelligent use of space and the intelligent use of language may form a mutually reinforcing pair, pursuing a common cognitive agenda.

Developmental investigations lend some substance to such a hypothesis. To take just one example, Namy *et al.* (1997) conducted a series of experiments involving children's use of space to represent similarity. Very briefly, what the experiments suggest is that spatial groupings of play objects (such as putting all the balls *here*, and all the boxes *there*) are not mere spatially expressed reflections of fully achieved grasp of category membership, but rather part and parcel of the process of coming to learn about categories and to discover the use of space as a means of representing category membership. The process the investigators document, in rich microgenetic detail, is one of bootstrapping that starts with early play experiences in which the child is interested in one kind of play object and hence ends up (as a side effect) with those objects grouped together in space. Such self-created groupings help the child to discover the possibility and the value of spatial classification itself. Crucial to this discovery is the child's engagement in preferential play in which one type of object is preferred over another. This kind of play was shown to lead, over relatively short periods of developmental time, to the emergence of true exhaustive classification behaviour, in which spatial organization functions as a symbolic indicator of category membership.

This whole process is one of incremental cognitive self-stimulation within a partially self-constructed cognitive niche. The perceptually available (grouped) products of the child's own activity form the new inputs that favour learning about exhaustive classification and (simultaneously) about the use of space as a means of representing category membership. The capacities of spontaneous spatial classification that this developmental bootstrapping helps create may then further scaffold the process of learning names and labels, while the acquisition of new names and labels in turn promotes the exploration of new and more sophisticated spatial groupings.

4. EPISTEMIC ENGINEERS

Our second example of cognitive niche construction emphasizes the transformative power of incrementally organized and actively engineered epistemic resources in the evolution and development of human cognition. To bring this phenomenon into focus, it helps to introduce the notion, due to Sterelny (2003), of *cumulative downstream epistemic engineering*. Sterelny offers an account of human uniqueness that gives pride of place to our extraordinary capacities as 'ecological engineers', that is to say, as the active constructors of our own cognitive niches. Having earlier argued for group selection as a key force in human evolution, Sterelny notes that groups of humans engineer their own habitats, and that these are transmitted to the next generation, who further modify the habitat. Importantly, some of these modifications are to the epistemic environment, and affect the informational structures and opportunities presented to each subsequent

generation. Although other animals clearly engage in niche construction, it is only in the human species (Sterelny argues) that we see this potent, cumulative, runaway (self-fuelling) process of epistemic engineering.

Niche construction is depicted by Sterelny as a kind of additional inheritance mechanism, working alongside (and interacting with) genetic inheritance. One of the points of interaction concerns phenotypic plasticity. For rampant niche construction yields a rapid succession of selective environments, and hence favours the (biological) evolution of phenotypic plasticity. Hominid minds, Sterelny suggests, are adapted to the spread of variation itself. To cope with such variability, we are said to have evolved powerful forms of developmental plasticity. These allow early learning to induce persisting and stable forms of neural reorganization, impacting our range of automatic skills, affective responses and generally reorganizing human cognition in deep and profound ways. The upshot is that 'the same initial set of developmental resources can differentiate into quite different final cognitive products' (Sterelny 2003, p. 166). In this way:

transforming hominid developmental environments transformed hominid brains themselves. As hominids remade their own worlds, they indirectly remade themselves.

(Sterelny 2003, p. 173)

We see this explanatory template in action in, for example, Sterelny's account of our capacity to interpret others as intentional agents. Thus:

Selection for interpretative skills could lead to a different evolutionary trajectory: selection on parents (and via group selection on the band as a whole) for actions which scaffold the development of the interpretative capacities. Selection rebuilds the epistemic environment to scaffold the development of those capacities.

(Sterelny 2003, p. 221)

Basic perceptual adaptations, for example, gaze monitoring, etc., are thus supposed to be bootstrapped up to a full-blown 'mind-reading' ability via the predictable effects of intense social scaffolding: the child is surrounded by exemplars of mind-reading in action; she is nudged by cultural inventions such as the use of simplified narratives⁶ (and, ultimately, books and pictures); prompted by parental rehearsal of her own intentions; and provided with a rich palate of linguistic tools such as words for mental states. Such 'incremental environmental engineering' provides, we are told, a 'wealth of the stimulus' argument against the innateness hypothesis (Sterelny 2003, p. 223). Our theory of mind, according to this argument, is not wired in at birth, but acquired by rich developmental immersion. Such immersion may itself have 'architectural consequences' (Sterelny 2003, p. 225), but these are the upshot, not the precondition, of learning. This explanatory strategy thus depicts much of what is most distinctive in human cognition as rooted in the reliable effects, on developmentally plastic brains, of immersion in a well-engineered, cumulatively constructed cognitive niche.

Sterelny's emphasis is thus very much upon the direct neural consequences of the culturally and

artefactually scaffolded training regimes applied to young human minds. But while such consequences are surely of the utmost importance, they do not yet exhaust the cognition-transforming effects of material artefacts and culture. For many of the new cognitive regimes supported by our best bouts of incremental epistemic engineering seem to resist full internalization. It is no use, as Ed Hutchins (personal communication) points out, trying to imagine a slide rule when you need to work out a log or cosine! Plastic human brains may nonetheless learn to factor the operation and information-bearing role of such external props and artefacts deep into their own problem-solving routines, creating hybrid cognitive circuits that are themselves the physical mechanisms underlying specific problem-solving performances. We thus come to our final and arguably most radical take on cognitive niche construction.

5. EXTENDED COGNITIVE SYSTEMS

Under certain conditions, non-organic props and aids, many of which are either culturally inherited tools or structures manipulated by culturally transmitted practices, might themselves count as proper parts of extended cognitive processes (e.g. Clark & Chalmers 1998; Hurley 1998; Rowlands 1999; Wilson 2004; Clark *in press*). Consider an accountant, Ada, who is extremely good at dealing with long tables of figures. Over the years, Ada has learnt how to solve specific classes of accounting problems by rapidly scanning the columns, copying some numbers onto a paper notepad, then looking to and from those numbers (carefully arrayed on the page) back to the columns of figures. This is all now second nature to Ada, who scribbles at lightning speed deploying a variety of ‘minimal memory strategies’ (Ballard *et al.* 1997). Instead of attempting to commit multiple complex numerical quantities and dependencies to biological short-term memory, Ada creates and follows trails through the scribbled numbers, relying on self-created external traces every time an intermediate result is obtained. These traces are visited and revisited on a ‘just in time, need to know’ basis, briefly shunting specific items of information into and out of short-term organic memory, in much the same way as a serial computer shifts information to and from the central registers in the course of carrying out some computation. This extended process may be best analysed as a set of problem-solving state transitions whose implementation happens to involve a distributed combination of organic memory, motor actions, external symbolic storage and just-in-time perceptual access.

Wilson’s (1994, 2004) notion of ‘wide computation’ captures the key features of such an extended approach. According to wide computationalism, ‘at least some of the computational systems that drive cognition reach beyond the limits of the organismic boundary’ (Wilson 2004, p. 165). The larger systems thus constituted are, Wilson insists, unified wholes such that ‘the resulting mind-world computational system itself, and not just the part of it inside the head, is genuinely cognitive’ (Wilson 2004, p. 167). Extended cognitive systems theorists thus reject the image of mind as a kind of input–output sandwich with cognition as the filling

(for this picture, and many more arguments for its rejection, see Hurley 1998; see also Clark & Chalmers 1998; Wheeler 2005). Instead, we confront an image of the local mechanisms of human cognition quite literally bleeding out into body and world.

6. DARWINIAN MODULES

And now for something completely different—or so it would seem. We have been mapping out an account of ourselves in which the human brain is depicted as a vortex of large-scale developmental and adaptive plasticity, positioned in an ongoing and co-determining interactive relationship with a dynamic flow of culturally evolving non-neural elements. However, what looks, on the face of things, to be a very different vision of our evolved neural engine and of how it relates to its cultural environment finds expression in the pages of the evolutionary psychology literature. It is time to scout that alternative vision.

Evolutionary psychology starts from the assumption that just as there are anatomical adaptations (bodily structures shaped by natural selection to solve certain adaptive problems), so there are psychological adaptations (internal information processing mechanisms shaped by natural selection to solve certain other adaptive problems). As Cosmides & Tooby (1987, p. 282) put it, ‘[the] evolutionary function of the human brain is to process information in ways that lead to adaptive behavior’. Evolutionary psychologists argue that it follows from this ‘Darwinized’ conception of information processing psychology that our innate cognitive endowment, as shared by all developmentally normal human beings, is not a domain-general learning and reasoning engine (as many social scientists and others have claimed), but rather (to use a now famous image) a psychological Swiss army knife, in that it comprises a large collection of specialized cognitive tools. This collection of tools is depicted as a suite of genetically specified, domain-specific computational mechanisms, often called *modules*, each of which (i) is triggered by informational inputs specific to a particular evolutionarily salient domain (e.g. choosing a mate and social exchange) and (ii) has access to internally stored information about that domain alone. Thus, the Swiss army knife account of mind is sometimes glossed as the massive modularity hypothesis (Sperber 1996; Samuels 1998).⁷

Two immediate clarifications of this picture are in order. First, it is important to note a distinguishing feature of the tabled approach to modularity. According to the evolutionary-psychological picture, the modules that comprise our innate cognitive endowment are to be demarcated at a functional level of analysis, an implication of which is that they need not be realized in localized regions of neural hardware (Gaulin & McBurney 2001). Secondly, evolutionary psychologists argue that in order to give an account of our adapted cognitive modules, one needs to identify the appropriate selective environment. This is a local application of a general principle. When one attempts to explain adaptation, one needs to have in view the ‘composite of environmental properties of the most recent segment of a species’ evolution that

encompasses the period during which its modern collection of adaptations assumed their present form' (Tooby & Cosmides 1990, p. 388). This crucial slice of selective history is what evolutionary psychologists call a trait's *environment of evolutionary adaptedness* (EEA). Of course, the relevant EEA may well not be the current environment in which a trait operates. Environments sometimes change, and evolution by cumulative Darwinian selection is typically thought of as a rather slow process that may lag well behind such change. This is especially probable in the case of a trait as complex as the human brain, embedded in an environment rich in historically unfolding cultural dynamics. Applying this logic, evolutionary psychologists typically argue that the last time any significant modifications were made by selection to the human brain's functional architecture was during the Pleistocene epoch (*ca* 2 Myr to 10 kyr ago), when humans were hunter-gatherers. So the composite of selection pressures at work in the Pleistocene constitutes our brain's EEA (see Crawford 1998 for discussion). This is where one finds the adaptive problems to which the modules housed by the modern brain—modules which have been inherited essentially unchanged from our Pleistocene hunter-gatherer ancestors—constitute evolved solutions.

Although the identification of the human EEA with the hunter-gatherer Pleistocene environment is an idea that has attracted a good deal of critical fire (e.g. Gould 2000; Smith *et al.* 2001), it does help the evolutionary psychologist to account for the fact that some of our behaviour fails to maximize fitness in modern cultural environments. For example, modern human males do not adopt a fitness-enhancing strategy of widespread sperm donation because our reproductive strategies are designed for Pleistocene conditions. And the fitness-decreasing obesity brought about by an overindulgence in sugar-rich foods in technologically advanced countries may be explained by the fact that our sweet tooth, which was adaptive in the nutritional challenges posed by the Pleistocene, has since been rendered maladaptive in such countries by the mass availability of refined sugar.

This image of a species-wide assemblage of evolved domain-specific information processing mechanisms, meshed with ancestral environmental factors, provides the background to a further aspect of the overall evolutionary-psychological picture that will be important in what follows. Evolutionary psychologists claim that behind all the manifest diversity in human cultural behaviour, there sits an *evolved universal human nature*. In what, then, does this evolved universal human nature consist, and how, given its alleged species-wide homogeneity, does it generate that remarkable diversity in cultural behaviour?

7. HUMAN NATURE

From what we have seen so far, it might seem that the evolutionary-psychological notion of an evolved universal human nature will be cashed out in terms of a suite of Darwinian modules possessed by all developmentally normal adult human beings. However, we need to be careful in how we handle this idea because

the fact is that that suite of modules, even as portrayed in evolutionary psychology, is not strictly universal. For example, whether or not a particular psychological adaptation is ultimately 'wired up' in a certain way in a specific individual will typically depend on the presence of certain environmental triggers that, under normal circumstances, occur reliably at critical stages during development. (For a dramatic example, consider the need for a rich linguistic environment to be present during language development.) Moreover, there may be alternative psychological adaptations available to development that are under the control of genetic switches (roughly, mechanisms by which genes are turned on or off through the absence or presence of DNA-binding proteins). Indeed, evolutionary psychologists argue that men and women confront divergent, sex-relative adaptive problems when it comes to finding, holding onto and reproducing with a mate. Thus, men and women instantiate different, sex-relative psychological adaptations in the mating game. Since sex determination is under the control of a genetic switch, so are these alternative psychological architectures.

What the existence of such alternative developmental trajectories demonstrates is that the suite of cognitive modules possessed by humankind is not strictly universal and so cannot constitute our species-wide human nature. What might then? The answer, nicely isolated by Buller (2005), is an evolved species-wide set of genetically specified developmental programs that (i) determine how the emerging human phenotype responds to critical environmental triggers and (ii) control processes such as genetic switching. It is at that level that strict universality (allegedly) holds, and at which our evolved human nature is (allegedly) to be found.

Now, if all developmentally normal human beings share a set of genetically specified developmental programs and, as a result, at least a very large number of innately specified psychological adaptations meshed with ancestral environments, what explains the variability of human behaviour across contemporary cultures? Here, we can draw a lesson from the example of ordinary digital computer programs. As in such programs, our cognitive information processing modules may respond differentially to variations in the inputs that they receive, inputs that are supplied largely by the particular cultural environments in which the bearers of those modules are embedded. A developmental version of this process is equally important. In certain cases, a particular innately specified module (e.g. a Chomskyan language acquisition device) may be exposed to different developmental environments (different linguistic communities providing different developmental inputs), leading ultimately to cognitive variation (different speakers learning and producing different languages).

Our second vision has now emerged fully. It is at root a vision of the evolved human brain as a locus of relatively static, genetically based forms of neural encoding and processing, executing a restricted set of pre-specified adapted functions in response to the triggers provided by variable cultural inputs. This certainly seems to suggest a very different view of what it is to be a natural human thinker from the one evoked by our synthesis of

embodied and extended cognition, cultural evolution and cognitive niche construction. But just how much of an intellectual chasm really exists between these apparently divergent views? In other words, along which dimensions, and to what extent, are our two visions in genuine competition with each other? It is to this issue that we shall now turn.

8. REMOULDING MODULARITY

What seems clear is that there is no necessary tension between, on one hand, an approach that foregrounds cultural evolution and, on the other, the kind of cognitive modularity favoured by the evolutionary psychologists. This might seem an odd claim to make at first, given that the fans of cultural evolution often place an emphasis on psychological mechanisms that exhibit a robust kind of domain generality. For example, drawing on Boyd & Richerson's (1985) dual inheritance model, a model that (as Sterelny's approach sketched earlier) stresses cultural as well as genetic transmission in evolution, Coultas (2004) provides experimental evidence that individual human beings have an essentially domain-general tendency to conform in social groups, a tendency that can be adaptive for the individual when information gathering by that individual would be costly. And Tomasello (1999), in a treatment that also stresses dual inheritance, argues that evolution has endowed us with a set of basic cognitive capacities, including shared attention and the imitation of other humans' behaviours and intentions, that allow us to take developmental advantage of a kind of accumulated species-specific knowledge made available through human cultural environments. At the heart of this process, and the capacity that sets human beings apart from other species, is our ability to identify intentions in others. It is this uniquely human, essentially domain-general ability, argues Tomasello, that allows us to build on foundational capacities that we share with other animals (such as the capacities for tool use and signalling), in order to become vastly more sophisticated thinkers in specific domains (e.g. vastly more sophisticated tool users and signallers) than have our evolutionary cousins. Finally, as we have seen already, Sterelny (2003) offers an account of our capacity to interpret others as intentional agents, according to which basic perceptual adaptations are bootstrapped up to a full-blown mind-reading ability via cognitive niche construction. This contrasts sharply with the evolutionary-psychological idea of an innate 'folk psychology' module, in the form of a domain-specific adaptation for mind-reading.

That said, Atran (2001, p. 8) presents an alternative view of the relationship between cultural transmission and cognitive modularity in which the latter underlies the former, with certain modules serving as 'as a principled basis for transmission and acquisition of more variable and extended forms of cultural knowledge'. For example, he argues that the widespread anthropological phenomenon of totemism—religious systems in which generic species spiritually represent social groups (e.g. an animal that spiritually represents a clan)—piggybacks on a genetically specified folk

biology module. That module latches onto generic species (and groups of generic species) whose intrinsically well-structured character renders them apt for memorability and cultural transmission between minds. These underlying categories supply cognitive hooks onto which our minds subsequently hang beliefs about intrinsically less well-structured social groups. In sum, according to Atran (2001, p. 8):

modularized structures—such as those which produce folkmechanical, folkpsychological and folkbiological concepts—are special players in cultural evolution. Their native stability derivatively attaches to more variable and difficult-to-learn representational forms, thus enhancing the latter's prospects for regularity and recurrence in transmission within and across cultures.

The availability of these alternative positions within the evolution-of-cognition research programme suggests strongly that one cannot infer that a cognitive architecture will be non-modular, or indeed that it will be modular, simply from the existence or otherwise of cultural transmission in the inheritance system.

Cognitive modularity is also compatible with the other partner in our symbiotic dyad, an embodied–extended approach to mind. A powerful illustration of how an embodied–extended modularity might go is provided by the field of *situated robotics* (e.g. Brooks 1991; Mataric 1991; Pfeifer & Bongard 2007). With the goal of building complete agents that are capable of integrating perception and action in real time so as to generate fast and fluid embodied adaptive behaviour, researchers in situated robotics shun the classical cognitive-scientific reliance on detailed internal world models, on the grounds that such structures are computationally expensive to build and keep up to date. Instead they adopt a design strategy according to which the robot regularly senses its environment to guide its action. It is this specific behaviour-generating strategy that marks out a robot as situated (Brooks 1991). Against this background, one of the key ideas from the field is that much of the richness and flexibility of intelligence is down not to general-purpose processes of reasoning and inference, but rather to integrated suites of special-purpose adaptive couplings that realize distributed or extended behaviour-generating strategies by combining non-trivial causal contributions from three constituencies: the brain (or its robotic equivalent); the non-neural body; and the environment. Moreover, this perspective provides one platform for the previously mentioned refusal to conceptualize perception and action as interfaces between mind and world. As Brooks (1991, p. 173) puts it, one of the guiding principles of the approach is that: 'There is no separation into perceptual system, central system, and actuation system. Pieces of the network [the distributed robotic control system] may perform more than one of these functions. More importantly, there is intimate intertwining of aspects of all three of them.'

A classic example of such work is provided by Maja Mataric's sonar-driven mobile robot, Toto (Mataric 1991). Toto wanders around its office environment following walls and avoiding obstacles. As it proceeds, it constructs an internal map based on landmarks,

which then enables it to navigate between locations. Toto is controlled by three main layers of situated special-purpose adaptive coupling: collision-free wandering; landmark detection; and map learning and path planning. What is theoretically interesting about Toto's map-learning and path-planning system is that navigation-related information is encoded in it in terms of patterns of embodied sensorimotor activity. For example, if, as Toto moves, it keeps detecting proximally located objects on its right-hand side, while its compass bearing remains unchanged, then a 'right wall' is encoded in its inner map, not as some agent-independent objectively specified entity, but in terms of its sensorimotor 'experience' at the time. These structured sensorimotor experiences (Toto's landmarks) are stored as connected nodes in a distributed graph, and this record of the robot's own embodied sensorimotor history constitutes its inner map of the spatial environment.

Crucially, given our interests, Toto's strategy of encoding spatial paths as internally represented sequences of past, current and expected embodied sensorimotor experiences is a domain-specific solution, one tailored to the particular navigational context for which the robot is designed. The action-oriented structures in question presumably would not be much good for a vast range of other space-related purposes, such as ordering correctly sized carpets for the corridors or determining the precise distance to the snack bar. Moreover, the navigation system is informationally encapsulated, in just the way required by the modularity hypothesis. (Of course, the map-learning and path-planning system depends on the successful functioning of the other layers of coupling, but informational encapsulation does not rule out such inter-systemic dependencies.) What all these suggest is that the sorts of situated special-purpose adaptive couplings promoted within situated robotics are illuminatingly understood as cognitive modules. Crucially, however, these modules have (what we might call) a horizontally extended character, in that their functional boundaries are no longer constrained by the orthodox transitions that remain in force in mainstream evolutionary psychology, between (i) perception and thought (in the world-to-body-to-mind input direction) and (ii) thought and action (in the mind-to-body-to-world output direction).

To develop further this notion of horizontally extended cognitive modularity, consider Ziemke *et al.*'s (2004) coevolutionary experiment involving two sets of simulated robots—scouts and drones—whose cooperation-demanding task is to enable the drones to find a spatially located goal. Both sets of agents are controlled by simple fixed topology neural networks under artificial evolutionary control. The task is posed in a grey-walled environment, in which each junction requiring a left turn to reach the goal is marked with a white stripe, while each junction requiring a right turn is marked with a black stripe. Scouts have cameras and so, in principle, can find their way to the goal autonomously using the turn-signalling stripes. By contrast, the drones have no cameras, only light sensors, so they cannot see the stripes. Their only hope, beyond random search, is to evolve to respond

correctly to light sources that are deposited by the scouts as the latter traverse the environment. So the scouts need to evolve a cognitive niche-construction strategy, one in which they place the light sources in such a way that they produce an increase in (what Ziemke *et al.* call) the *cognitive congeniality* of the environment inherited by the drones.⁸

Under the experimental conditions described, scouts evolve to drop light sources in response to the white stripe on the wall—thereby constructing a niche that simplifies the problem task for the drones—and drones evolve to exploit these 'road signs', by turning left while sensing the light, but right at the other junctions. This niche-construction scenario once again displays the distinctive hallmarks of situated special-purpose adaptive coupling (e.g. tight linkages between particular embodied sensorimotor capacities and task-dedicated action-generating strategies that factor in the reliable presence of specific environmental structures) and thereby of horizontally extended modularity.⁹

In spite of these positive steps towards a reconciliation between an approach that emphasizes cognitive modularity and one that emphasizes cultural transmission and the embodied–extended mind, an important issue remains to be addressed. As we have seen, evolutionary psychologists explain the development of cognitive modules in terms of a species-wide set of genetically specified developmental programs that orchestrate the journey from genotype to phenotype, and in particular from genes to massive modularity. This genocentric stance might seem to clash unhelpfully with an account of development that routinely appeals to the bootstrapping up of basic capacities via cultural transmission and cognitive niche construction, and which thereby shifts the centre of explanatory gravity away from genetic specification and towards a distributed matrix of co-determining genetic and environmental factors. Even here, however, there is some hope that the tension may be relieved, if we combine the thought that progressive modularization may emerge during development and learning (e.g. Karmiloff-Smith 1992), with an account of the conditions under which, within the sort of distributed developmental matrix just highlighted, genes may rightly be said to code for phenotypic traits (Wheeler & Clark 1999; Wheeler 2003). Each of these ideas warrants discussion.

9. EMERGENT MODULARITY

Karmiloff-Smith (1992) provides a compelling account of how, given the plasticity of early neural development, a progressive functional modularization may be realized by the brain as part of the developmental process. Evidence from cases of early brain damage indicate a degree of baseline neural plasticity that goes well beyond that suggested by the evolutionary-psychological image of a set of genetically specified modules installed in response to environmental triggers. The mind is not pre-structured at birth to be modular. Instead, a process of modularization is kick-started by a limited range of multilevel domain-specific predispositions that focus the young infant's attention

on certain proprietary inputs. The progressive development of emergent modular structures then proceeds interactively as these proprietary inputs in turn affect the development of the brain.

A rich example of how functional modularization may be the outcome of constrained dynamic interaction during development is provided by Hirsh-Pasek & Golinkoff's (1996) three-phase coalition model of language comprehension (see also Hollich *et al.* 2000). According to this model, infants in the first phase build on rudimentary language comprehension achieved during the second half of the first year of life to perform an initial segmentation of the flux of their acoustic and visual environments. On the basis of dispositions to note certain acoustic and visual cues, alongside a capacity for distributional and correlational analysis across phonological and rhythmic patterns of speech, the infant's task is to parcel up the flow of speech around her into acoustic units that will later become linguistically relevant, and to use these acoustic units to help her uncover highly significant structures in her environment (e.g. important events and objects). The second phase involves the interpretation of the acoustic units as components that correlate with linguistic categories (such as subject, verb and object), plus the mapping of individual word units onto their referents. In this way, semantics dislodges sound as the primary regulator of emerging language comprehension. Although during this phase children are beginning to comprehend multiword sentences and the role of word order in determining grammatical relations, such advanced comprehension is fragile, in that it depends on all the relevant social, semantic and syntactic cues being present. Thus, a supporting coalition of environmental factors forms a developmentally crucial cognitive scaffold. In the third phase, this dependency is overcome. The child's syntactic system becomes fully established, as indicated by the late onset ability to understand linguistic constructions that violate word-order assumptions (e.g. the English passive).

For Hirsh-Pasek & Golinkoff, then, language comprehension is kick-started by a system that is primed with dispositions to note salient inputs and their likelihood of occurring together. Functional modularization, in the form of a domain-specific, informationally encapsulated system for language comprehension, develops progressively through interactive environmental engagement. So cognitive modularity may result from distributed developmental bootstrapping that potentially involves cultural transmission and cognitive niche construction. There seems to be no reason to think that there could not be a large number of such modules, so in that sense at least the human cognitive system may be a locus of massive emergent modularity—an emergent cognitive Swiss army knife! But now what about the evolutionary-psychological claim that cognitive modules are genetically specified? One might think of this as a key component of the evolutionary-psychological vision. What remains of this claim in the alternative story? The answer, we suggest, is: rather more than you might expect.

10. GENES, CODES AND EXPLANATORY SPREAD

There is a generic phenomenon that the present authors once dubbed explanatory spread (Wheeler & Clark 1999). Mameli (2005, p. 388) gives a clear exposition of what it entails.

Causal spread occurs when we discover some new factor causally involved in the occurrence of a phenomenon. Explanatory spread occurs when we realize that some factor that was not considered to be necessary in the explanation of a phenomenon is instead explanatorily necessary for that phenomenon. Or, to put it differently, explanatory spread occurs when we realize that some factor that was not taken to be part of a sufficient explanation of a phenomenon needs to be included in such explanation. Since the fact that something is causally required does not entail that it is also explanatorily required, causal spread does not necessarily lead to explanatory spread. But in cases where the newly discovered causal factor is deemed to be an important one, causal spread is likely to generate the inclusion of the newly discovered factor in any sufficient explanation of a phenomenon to which this factor causally contributes. That is, in these cases, causal spread leads to explanatory spread.

Where the phenomenon of interest is phenotypic form, the received position is that such structure is down to genetic specification. So one would have explanatory spread where one discovered a distributed developmental system in which non-genetic organismic and/or wider environmental factors made explanatorily non-negligible contributions to phenotypic form. That is the general picture on offer from approaches that emphasize cultural evolution, cognitive niche construction and (we can now add) emergent modularity.

So what? Crucially, some authors have argued that a proper recognition of developmental explanatory spread should lead us to reject the claim that genes specify phenotypic traits. Cognitive modules are, of course, examples of phenotypic traits, so if this anti-specification argument is sound, it would undermine the claim that such modules are genetically specified, and so re-establish a conflict between our two visions. But is that argument sound? To answer that question, let us consider a specific statement of it:

We have often heard it said that genes contain the 'information' that specifies a living being... [but] when we say that DNA contains what is necessary to specify a living being, we divest these components... of their interrelation with the rest of the network. It is the network of interactions in its entirety that constitutes and specifies the characteristics of a particular cell, and not one of its components. That modifications in the components called genes dramatically affect the structure is very certain. The error lies in confusing essential participation with unique responsibility. By the same token one could say that the political constitution of a country determines its history. This is obviously absurd. The political constitution is an essential component in any history but it does not contain the 'information' that specifies that history.

(Maturana & Varela 1987, p. 69)

What is going on here? The first thing to note (as the opening sentence of the above passage indicates) is that to conceive of genes as trait specifiers is to conceive of genes as developmental information carriers, i.e. as *coding for* phenotypic traits. Thus, much here turns on how one understands the nature of that coding relationship. It seems to us that Maturana and Varela's argument depends implicitly on a deceptively tempting, but ultimately flawed, view of coding talk that we call *strong instructionism* (Wheeler & Clark 1999; see also Wheeler 2003, 2006). Strong instructionism is the claim that what it means for some element to code for an outcome is for that element to *fully specify* the distinctive features of that outcome, where 'full specification' requires that those distinctive features may be predicted purely on the basis of what may be known about the putatively coding factor. In the present context, strong instructionism amounts to the claim that what it means for a gene (or a complex of genes) to code for a phenotypic trait is for that gene (or complex of genes) to fully specify the form of that trait. It is this kind of picture that is seemingly suggested by the classic Lorenzian image of the non-genetic material causes in development as the bricks and mortar out of which the organism is assembled according to a genetic blueprint (Lorenz 1965). However, given the presence of developmental explanatory spread (what Maturana and Varela call 'the network of interactions in its entirety'), the fact is that knowing the entire sequence of an organism's DNA will *not* be sufficient to predict phenotypic form. It is this point that underwrites Maturana and Varela's observation that the fan of genetic information mistakenly confuses 'essential participation with unique responsibility'. So, if the understanding of genes as coding for phenotypic traits is tied to strong instructionism, then, given developmental explanatory spread, that understanding is false.¹⁰

The trick, then, is to free coding talk about genes from strong instructionism. Fortunately, there is plenty of evidence that coding talk in other domains does not impose the full-specification condition. Indeed, in familiar cases of algorithms, programs, instruction sets and other such coding elements, those states and processes are able to perform their outcome-generating functions only given some assumed backdrop of other causally active states and processes (e.g. working operating systems) that themselves bear some of the responsibility for the exact form of the outcome produced. In other words, strong instructionism is a spectre without much of a haunting pedigree. That said, a word of warning: we need to avoid falling into the opposite trap of giving an account of genetic coding so excessively liberal, that where explanatory spread is present, *too many* developmental factors qualify as coding for phenotypic outcomes. For then the claim that a certain gene (or complex of genes) codes for some trait will simply fail to single out that gene (or complex of genes) as performing a distinctive developmental function.¹¹

To take just one example (for several others, see Wheeler 2006), say we adopted the superficially attractive view that genes code for traits insofar as they are what is passed on from one generation to the next in evolution. If we define inheritance without an

antecedent pro-gene prejudice, as the biological like-begets-like phenomenon, and so as to fix on elements that are robustly and reliably replicated in each generation of a lineage, and that persist long enough to be the target of cumulative selection, then the fact seems to be that genes are *not* all that organisms inherit. For example, there are the so-called *epigenetic inheritance systems*, such as the inheritance of methylation patterns via a separate (i.e. from the genetic) copying system; and there is inheritance through *host imprinting*, as when parasitic birds, born in the nest of a host species, imprint on that nest as chicks, and then later lay their own eggs in the nest of that species; and then there is inheritance via our old friend *niche construction*, as when beaver offspring inherit both the dam that was communally constructed by the previous generation and the altered river flow that that physical structure has produced. What this indicates is that if being inherited is sufficient for some developmental factor to qualify as coding for a phenotypic trait, then non-genetic factors will regularly count as coding elements, which violates our excessive liberality constraint.

There is, of course, much more to be said about this issue. In the present treatment, we have done little more than sketch the form that an account of coding talk would have to take, if it is to allow genes to code for (and thus, in a robust sense, specify) phenotypic traits (including cognitive modules), even in the midst of an explanatory spread that involved cultural transmission and cognitive niche construction. But, if we can successfully navigate between the Scylla of strong instructionism and the Charybdis of excessive liberality, we would potentially have access to such an account. Allied with the concept of emergent modularity, that result would do much to effect a rapprochement between our alternative visions of evolved human cognition.¹²

11. HUMAN NATURE RECONSIDERED

It is time to revisit the evolutionary psychologist's notion of an evolved universal human nature, conceived as a species-wide set of genetically specified developmental programs that orchestrate the journey from genotype to phenotype. According to this view, a maturing human being, embedded in a normal developmental environment, will end up with a particular, species-wide set of cognitive modules (allowing for some branching pathways, e.g. between the sexes). Significant challenges to this view are posed by the powerful role assigned, by cognitive niche-construction models, to stacked sequences of training environments in the emergence of specific functional modules.¹³ While the early stages of such key developmental trajectories may, as we saw, be rather predictably determined by small native biases, the later stages often reflect both the cumulative effects of cultural evolution and transmission, and the potent effects of the ongoing self-selection of training environments. A child whose early experience is shaped by the special environments provided by books and software programs, and whose own emerging cognitive profile favours certain elements within that culturally enabled nexus over other elements, will end up with a cognitive system that is not just superficially, but profoundly,

different from that of a differently encultured child. Such a view finds expression in, for example, Schlesinger & Parisi's (2007, p. 153) notion of an emergent constraint according to which:

the outcome of a developmental process need not be programmed in by maturation but instead may occur as the result of successive learning experiences that the organism determines or selects for itself.

The neuroroboticist Olaf Sporns describes the larger situation well, noting that:

[the] architecture of the brain... and the statistics of the environment [are] not fixed. Rather, brain-connectivity is subject to a broad spectrum of input-, experience-, and activity-dependent processes which shape and structure its patterning and strengths (Johnson 2001). These changes, in turn, result in altered interactions with the environment, exerting causal influences on what is experienced and sensed in the future.

(Sporns 2007, p. 179)

This kind of 'neuroconstructivist' framework (for a compelling array of worked examples, see Mareschal *et al.* 2007a,b) helps locate a potential challenge for any notion of an evolved human nature that ties that nature too closely to the properties and features of the EEA. For what is special about human brains, and what best explains the distinctive features of human intelligence, may be precisely their ability (courtesy of extended development and extensive neural plasticity) to enter into deep, complex and ultimately architecture-determining relationships with an open-ended variety of culturally transmitted practices, endowments and non-biological constructs, props and aids. Perhaps it is because our brains, more than those of any other animal on the planet, are primed to seek and consummate such intimate relations with non-biological resources that we end up as bright and as capable of abstract thought as we are. If so, our distinctive universal human nature, insofar as it exists at all, would rather be a nature of biologically determined openness to deep, learning- and development-mediated, change.

It is at this point that we locate a potential challenge to the evolutionary psychologists' specific vision of a universal human nature. For that vision, as we saw earlier, commits them to a restricted range of potential cognitive modules, with that range determined by a suite of genetically specified developmental programs. As a result, the range of possible normal variation among cognitive modules is strictly and endogenously limited. By contrast, the constructivist vision of horizontally extended and emergent cognitive modules places no such clean limits upon the range of variation. Insofar as there is something worth calling a universal human nature on this alternative view, that nature lies precisely in our continual openness to radical cognitive change. Our fixed nature is thus a kind of meta-nature: the suite of capacities, practices and proclivities that enable the development, use and propagation of a much more open-ended set of horizontally extended and emergent cognitive modules.

Such openness, as stressed by recent works on embodied and extended cognition, adds important complexity to accounts that emphasize the EEA. For

we must now take into account a plastic evolutionary overlay that yields a constantly moving target, an extended cognitive architecture whose constancy lies mainly in its continual openness to change. Even granting that the biological innovations that got this ball rolling may have consisted only in some small tweaks to an ancestral repertoire, the upshot of this subtle alteration would be a sudden, massive leap in cognitive-architectural space: the emergence of a cognitive machine intrinsically geared to self-transformation, artefact-based expansion and a snowballing/bootstrapping process of computational and representational growth. The machinery of human reason (the environmentally extended apparatus of our distinctively human intelligence) could thus turn out to be rooted in a biologically incremental progression while simultaneously existing on the far side of a precipitous cliff in cognitive-architectural space.

12. CONCLUSIONS: THE SPACE BETWEEN

Such, at least, would be the most radical model, one that indeed locates some genuine tension between the evolutionary psychologist's emphasis on hard modules and the EEA, and the cognitive niche constructivist emphasis on emergent modularity as reflecting the complex ratchet effects made available by the interplay of neural plasticity, learning and embodied activity involving inherited or self-created environmental structure.

But between these poles of human nature as highly reflective of the specific features of the EEA, and human nature as one of extensive openness to training and input-based modification, lies the full and inviting cognitive space structured by the triple helix of culture, embodiment and genes. Triple helix models of mind recognize the role of genetic biases in sculpting key developmental trajectories, and the resulting space both for strong forms of genetically specified cognitive modularity and for weaker forms of emergent modularity resulting from trajectories marked by multiple bouts of culturally scaffolded experience and the self-selection of environments. But crucially, the triple helix template also invites us to consider, pretty much on a case-by-case basis, all points and stations in between. Understanding this spectrum, and unravelling the complex interplay between genes, environments and embodied action, will surely be one of the great intellectual adventures of the 21st century.

This paper was prepared in part thanks to support granted to A.C. by the AHRC, under the ESF Eurocores CNCC scheme, as part of the CONTACT (Consciousness in Interaction) project AH/E511139/1, and to M.W. by the AHRC as part of project AH/F002963/1. Some sections have been adapted from Clark (2003, in press, ch. 4) and Wheeler (2006, in press). Many thanks to John Protevi, Kenny Smith and John Sutton for their constructive critical feedback on an earlier version of this paper.

ENDNOTES

¹In line with much contemporary usage, we shall take the term 'evolutionary psychology' to signal not simply any psychological science that takes its cues from evolutionary biology, but rather a specific research paradigm centred on the work of Cosmides & Tooby (1987), Buss (1994) and Pinker (1997), among others.

²The idea of a triple helix in evolution was originally developed by Richard Lewontin (2000), who identified its components as genes, organism and environment. Our usage makes contact with Lewontin's own, but adapts the latter two components so as to focus on the especially potent and intriguing dimensions provided by embodiment and culture.

³For a host of other examples, see Laland *et al.* (2000) and Odling-Smee *et al.* (2003). See also Dawkins (1982), Lewontin (1983) and Turner (2000).

⁴The basic idea of human beings as cognitive niche constructors is familiar within cognitive science. Richard Gregory (1981) spoke of 'cognition amplifiers', Don Norman (1993) of 'things that make us smart', Kirsh & Maglio (1994) of 'epistemic actions' and Daniel Dennett (1996) of 'tools for thought'.

⁵It is worth noting that nothing in this view commits us to the notion of a single 'abstract' human subject rather than a population of subjects with different traits and nuances. Instead, it is best to think of a range of subjects displaying, as a result of genetic, cultural and environmental influences, a spread of different traits and capacities. For each such trait and capacity, taken in its local context, there will be a correlated pattern of empowerment and constraint. The most successful human groups will then be those in which the spread itself (which will include differences in affect and affective response) is mutually beneficial. Thanks to John Protevi (personal communication) for drawing these issues to our attention.

⁶For a compelling analysis of how involvement in a particular kind of narrative practice may explain the developmental path to an understanding of other minds, an understanding which itself turns on the construction of narratives, see Hutto (2008).

⁷Here, we shall not be concerned with assessing the positive conceptual arguments or the experimental data that are supposed to take us from the Darwinization of information processing psychology to the massive modularity of the adapted mind. In general, the conceptual arguments turn on the thought that domain-general mechanisms in isolation, i.e. without assistance from domain-specific mechanisms, would not be able to solve the adaptive problems confronted by the brain, or at least that any domain-general mechanism in the evolving population will typically have been systematically outperformed by any competing domain-specific mechanisms, such that it is the latter kind of mechanism that will have been selected for. For critical discussion of the arguments and evidence here that typically finds them wanting, see, for example, Samuels (1998), Sterelny & Griffiths (1999), Atkinson & Wheeler (2004) and Buller (2005).

⁸The fitness scores that determine the survival and reproduction prospects in the evolutionary scenario are calculated as follows: in each trial, an individual scout is rewarded (i) for finding the goal itself and (ii) if an associated follower drone also reaches the goal, while each drone is rewarded (iii) for finding the goal itself and (iv) if an associated leader scout also reaches the goal. Thus, it is the achievement of the goal state that is rewarded directly, and not the specific strategies for reaching that state.

⁹One might see the drones as constituting a limit case in which the entire control system implements a single functionally identified module.

¹⁰It might seem that strong instructionism about genes is a straw position that no one seriously holds. However, the fact is that the idea remains insidiously at work behind commonplace metaphors for understanding the relationship between genes and traits. As John Dupre, a philosopher of biology and the director of a centre for research on genomics in society (*Egenis*, <http://www.genomicsnetwork.ac.uk/egenis>), comments: 'It is still common to hear the genome described, for instance, even by eminent experts, as a blueprint for the organism... Perhaps not many people will defend the blueprint metaphor very far these days, if pushed, however. A common retreat is to the metaphor of a recipe... But this metaphor is still quite inadequate. With due allowance for an element of assumed common knowledge, the recipe is a complete set of instructions for how to make the cake.' (Dupre 2005, p. 198, our emphasis)

¹¹A longer justification for why such liberality is excessive goes like this. If the primary goal of introducing the concept of genetic coding is to single out genes as privileged causal elements in the developmental process, then it might well seem that any successful account of coding talk must have the consequence that, of the many causal factors that

combine causally during development, it is the genes *alone* that end up coding for phenotypic traits. Elsewhere one of us has dubbed this *the uniqueness constraint* (Wheeler 2006). Griffiths & Knight (1998; see also Griffiths 2001) introduce what is essentially the same constraint in terms of what they call the 'parity thesis'. The uniqueness constraint will not be met if either (i) the account of genetic coding under consideration fails to deliver the result that genes code for traits, since if genes do not code for traits then they can not do so uniquely, or (ii) that account does deliver the result that genes code for traits, but its conditions for what it is to do this are met by other elements in the extended developmental system, since then genes will not be the only developmental elements that code for traits. Condition (ii) gives expression to the excessive liberality problem. For discussion and a more careful formulation of the uniqueness constraint, see Wheeler (2006).

¹²Note that the extended character of certain embodied and situated modules is no barrier to this project. Dawkins' (1982) influential notion of the extended phenotype already shows us how genes may be understood as coding for traits that are located outside the skin of the organism (e.g. the genes that code for the spider's web). Beyond that, however, the waters between our sea monsters are exceptionally turbulent. For example, the present authors have argued in the past (Wheeler & Clark 1999; Wheeler 2003) that what we need is an account of genetic coding based on two features of protein synthesis: the arbitrariness of the mappings from particular nucleotide triplets to particular amino acids, and the way in which information is consumed by the subsystems that implement translation. However, one of us (Wheeler 2006) has subsequently argued that once the details of this account are filled in, it turns out that, strictly speaking, it is not the molecules of DNA that code in development, but rather the downstream nucleotide triplets out of which molecules of mRNA are constructed. It may be that the final route between the dual dangers of strong instructionism and excessive liberality is still to be found.

¹³For a different way of criticizing the evolutionary-psychological conception of human nature, one that identifies an alleged inconsistency between that conception and the population-thinking foundations of contemporary neo-Darwinian biology, see Buller (2005).

REFERENCES

- Atkinson, A. & Wheeler, M. 2004 The grain of domains: the evolutionary-psychological case against domain-general cognition. *Mind Lang.* **19**, 147–176. (doi:10.1111/j.1468-0017.2004.00252.x)
- Atran, S. 2001 The case for modularity: sin or salvation? *Evol. Cogn.* **7**, 46–55.
- Ballard, D., Hayhoe, M., Pook, P. & Rao, R. 1997 Deictic codes for the embodiment of cognition. *Behav. Brain Sci.* **20**, 723–767.
- Barkow, J. H., Cosmides, L. & Tooby, J. (eds) 1992 *The adapted mind: evolutionary psychology and the generation of culture*. New York, NY: Oxford University Press.
- Beach, K. 1988 The role of external mnemonic symbols in acquiring an occupation. In *Practical aspects of memory* (eds M. M. Gruneberg & R. N. Sykes), pp. 342–346. New York, NY: Wiley.
- Boyd, R. & Richerson, P. J. 1985 *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.
- Brooks, R. A. 1991 Intelligence without reason. In *Proc. 12th Int. Joint Conf. on Artificial Intelligence*, pp. 569–595. San Mateo, California: Morgan Kaufman.
- Buller, D. J. 2005 *Adapting minds: evolutionary psychology and the persistent quest for human nature*. Cambridge, MA: MIT Press.
- Buss, D. M. 1994 *The evolution of desire: strategies of human mating*. New York, NY: Basic Books.
- Buss, D. M. 2005 *The handbook of evolutionary psychology*. Hoboken, NJ: Wiley.
- Clark, A. 1997 *Being there: putting brain, body and world together again*. Cambridge, MA: MIT Press.

- Clark, A. 2001 *Mindware: an introduction to the philosophy of cognitive science*. New York, NY: Oxford University Press.
- Clark, A. In press. *Supersizing the mind: embodiment, action, and cognitive extension*. New York, NY: Oxford University Press.
- Clark, A. & Chalmers, D. 1998 The extended mind. *Analysis* **58**, 7–19. (doi:10.1111/1467-8284.00096)
- Collins, S. H., Ruina, A. L., Tedrake, R. & Wisse, M. 2005 Efficient bipedal robots based on passive-dynamic walkers. *Science* **307**, 1082–1085. (doi:10.1126/science.1107799)
- Cosmides, L. & Tooby, J. 1987 From evolution to behavior: evolutionary psychology as the missing link. In *The latest on the best: essays on evolution and optimality* (ed. J. Dupre), pp. 227–306. Cambridge, MA: MIT Press.
- Coultas, J. C. 2004 When in Rome... an evolutionary perspective on conformity. *Group Processes and Intergroup Relations* **7**, 317–331. (doi:10.1177/1368430204046141)
- Crawford, C. 1998 Environments and adaptations: then and now. In *Handbook of evolutionary psychology: ideas, issues, and applications* (eds C. Crawford & D. L. Krebs), pp. 275–302. Mahwah, NJ: Lawrence Erlbaum.
- Dawkins, R. 1982 *The extended phenotype*. Oxford, UK: Oxford University Press.
- Dennett, D. C. 1996 *Kinds of minds: towards an understanding of consciousness*. London, UK: Weidenfeld & Nicolson.
- Dupre, J. 2005 Are there genes? In *Philosophy, biology and life* (ed. A. O'Hear), pp. 193–210. Cambridge, MA: Cambridge University Press.
- Feldman, M. W. & Cavalli-Sforza, L. L. 1989 On the theory of evolution under genetic and cultural transmission with application to the lactose absorption problem. In *Mathematical evolutionary theory* (ed. M. W. Feldman), pp. 145–173. Princeton, NJ: Princeton University Press.
- Gaulin, S. J. C. & McBurney, D. H. 2001 *Psychology: an evolutionary approach*. Upper Saddle River, NJ: Prentice Hall.
- Gould, S. J. 2000 More things in heaven and earth. In *Alas poor Darwin: arguments against evolutionary psychology* (eds H. Rose & S. Rose), pp. 101–126. New York, NY: Harmony Books.
- Gregory, R. 1981 *Mind in science: a history of explanations in psychology*. Cambridge, MA: Cambridge University Press.
- Griffiths, P. E. 2001 Genetic information: a metaphor in search of a theory. *Philos. Sci.* **68**, 394–412. (doi:10.1086/392891)
- Griffiths, P. E. & Knight, R. D. 1998 What is the developmentalist challenge? *Philos. Sci.* **65**, 253–258. (doi:10.1086/392636)
- Hirsh-Pasek, K. & Golinkoff, R. M. 1996 *The origins of grammar: evidence from early language comprehension*. Cambridge, MA: MIT Press.
- Hollich, G., Hirsh-Pasek, K., Tucker, M. L. & Golinkoff, R. M. 2000 A change is afoot: emergentist thinking in language acquisition. In *Downward causation* (eds P. Anderson, C. Emmeche, N. O. Finnemann & P. V. Christiansen), pp. 143–178. Oxford, UK: Aarhus University Press.
- Hurley, S. L. 1998 *Consciousness in action*. Cambridge, MA: Harvard University Press.
- Hutto, D. 2008 *Folk psychological narratives: the sociocultural basis of understanding reasons*. Cambridge, MA: MIT Press.
- Jablonka, E. & Lamb, M. J. 2005 *Evolution in four dimensions: genetic epigenetic, behavioral, and symbolic variation in the history of life*. Cambridge, MA: MIT Press.
- Johnson, M. 2001 Functional brain development in humans. *Nat. Rev. Neurosci.* **2**, 475–483. (doi:10.1038/35081509)
- Kamiloff-Smith, A. 1992 *Beyond modularity: a developmental perspective on cognitive science*. Cambridge, MA: MIT Press.
- Kirby, S. 2002 Learning, bottlenecks and the evolution of recursive syntax. In *Linguistic evolution through language acquisition: formal and computational models* (ed. E. Briscoe), pp. 173–204. Cambridge, MA: Cambridge University Press.
- Kirsh, D. 1995 The intelligent use of space. *Artif. Intell.* **73**, 31–68. (doi:10.1016/0004-3702(94)00017-U)
- Kirsh, D. & Maglio, P. 1994 On distinguishing epistemic from pragmatic action. *Cogn. Sci.* **18**, 513–549. (doi:10.1016/0364-0213(94)90007-8)
- Laland, K. N., Odling-Smee, J. & Feldman, M. W. 2000 Niche construction, biological evolution and cultural change. *Behav. Brain Sci.* **23**, 131–146. (doi:10.1017/S0140525X00002417)
- Lewontin, R. C. 1983 Gene, organism, and environment. In *Evolution from molecules to men* (ed. D. S. Bendall), pp. 273–285. Cambridge, MA: Cambridge University Press.
- Lewontin, R. C. 2000 *The triple helix: gene, organism and environment*. Cambridge, MA: Harvard University Press.
- Lorenz, K. 1965 *Evolution and the modification of behaviour*. Chicago, IL: University of Chicago Press.
- Mameli, M. 2005 The inheritance of features. *Biol. Philos.* **20**, 365–399. (doi:10.1007/s10539-004-0560-0)
- Mareschal, D., Johnson, M., Sirois, S., Spratling, M., Thomas, M. & Westermann, G. 2007a *Neuroconstructivism: vol. 1, how the brain constructs cognition*. New York, NY: Oxford University Press.
- Mareschal, D., Sirois, S., Westermann, G. & Johnson, M. 2007b *Neuroconstructivism: vol. 2, perspectives and prospects*. New York, NY: Oxford University Press.
- Mataric, M. 1991 Navigating with a rat brain: a neurobiologically inspired model for robot spatial representation. In *From animals to animats: proceedings of the first international conference on simulation of adaptive behavior* (eds J.-A. Meyer & S. Wilson), pp. 169–175. Cambridge, MA: MIT Press.
- Maturana, H. & Varela, F. J. 1987 *The tree of knowledge: the biological roots of human understanding*. Boston, MA: New Science Library.
- Namy, L., Smith, L. & Gershkoff-Stowe, L. 1997 Young children's discovery of spatial classification. *Cogn. Dev.* **12**, 163–184. (doi:10.1016/S0885-2014(97)90011-3)
- Noë, A. 2004 *Action in perception*. Cambridge, MA: MIT Press.
- Norman, D. 1993 *Things that make us smart*. Cambridge, MA: Perseus Books.
- Odling-Smee, J., Laland, K. & Feldman, M. 2003 *Niche construction*. Princeton, NJ: Princeton University Press.
- Pfeifer, R. & Bongard, J. 2007 *How the body shapes the way we think*. Cambridge, MA: MIT Press.
- Pinker, S. 1997 *How the mind works*. New York, NY: Norton.
- Rowlands, M. 1999 *The body in mind*. Cambridge, MA: Cambridge University Press.
- Samuels, R. 1998 Evolutionary psychology and the massive modularity hypothesis. *Br. J. Philos. Sci.* **49**, 575–602. (doi:10.1093/bjps/49.4.575)
- Schlesinger, M. & Parisi, D. 2007 Connectionism in an artificial life perspective: simulating motor, cognitive, and language development. In *Neuroconstructivism: perspectives and prospects* (eds D. Mareschal, S. Sirois, G. Westermann & M. H. Johnson), pp. 129–158. Oxford, UK: Oxford University Press.
- Smith, E. A., Borgerhoff Mulder, M. & Hill, J. 2001 Controversies in the evolutionary social sciences: a guide for the perplexed. *Trends Ecol. Evol.* **16**, 128–135. (doi:10.1016/S0169-5347(00)02077-2)
- Sperber, D. 1996 *Explaining culture: a naturalistic approach*. Oxford, UK: Blackwell.

- Sporns, O. 2007 What neuro-robotic models can teach us about neural and cognitive development. In *Neuroconstructivism: perspectives and prospects* (eds D. Mareschal, S. Sirois, G. Westermann & M. H. Johnson), pp. 179–204. Oxford, UK: Oxford University Press.
- Sterelny, K. 2003 *Thought in a hostile world: the evolution of human cognition*. Oxford, UK: Blackwell.
- Sterelny, K. & Griffiths, P. E. 1999 *Sex and death: an introduction to philosophy of biology*. Chicago, IL: University of Chicago Press.
- Tomasello, M. 1999 *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
- Tooby, J. & Cosmides, L. 1990 The past explains the present: emotional adaptations and the structure of ancestral environments. *Ethol. Sociobiol.* **11**, 375–424. (doi:10.1016/0162-3095(90)90017-Z)
- Turner, S. J. 2000 *The extended organism: the physiology of animal-built structures*. Cambridge, MA: Harvard University Press.
- Varela, F. J., Thompson, E. & Rosch, E. 1991 *The embodied mind: cognitive science and human experience*. Cambridge, MA: MIT Press.
- von Frisch, K. 1975 *Animal architecture*. London, UK: Hutchinson.
- Wheeler, M. 2003 Do genes code for traits? In *Philosophical dimensions of logic and science: selected contributed papers from the 11th International Congress of Logic, Methodology, and Philosophy of Science* (eds A. Rojszczak, J. Cachro & G. Kurczewski), pp. 151–164. Dordrecht, The Netherlands: Kluwer.
- Wheeler, M. 2005 *Reconstructing the cognitive world: the next step*. Cambridge, MA: MIT Press.
- Wheeler, M. 2006 Traits, genes and coding. In *Handbook of the philosophy of biology*, pp. 381–411. Amsterdam, The Netherlands: Elsevier.
- Wheeler, M. In press. Evolutionary models in psychology. In *The Routledge companion to the philosophy of psychology* (eds P. Calvo and J. Symons). Oxford, UK: Routledge.
- Wheeler, M. & Clark, A. 1999 Genic representation: reconciling content and causal complexity. *Br. J. Philos. Sci.* **50**, 103–135. (doi:10.1093/bjps/50.1.103)
- Wilson, R. A. 1994 Wide computationalism. *Mind* **103**, 351–372. (doi:10.1093/mind/103.411.351)
- Wilson, R. A. 2004 *Boundaries of the mind: the individual in the fragile sciences*. Cambridge, MA: Cambridge University Press.
- Wilson, R. & Clark, A. In press. How to situate cognition: letting nature take its course. In *Cambridge handbook of situated cognition* (eds M. Aydede and P. Robbins). Cambridge, MA: Cambridge University Press.
- Ziemke, T., Bergfeldt, N., Buason, G., Susi, T. & Svensson, H. 2004 Evolving cognitive scaffolding and environment adaptation: a new research direction for evolutionary robotics. *Connect. Sci.* **16**, 339–350. (doi:10.1080/09540090412331314821)